

AFRL-SN-WP-TP-2003-100

**sarMAPPER: A REAL-TIME
INTERACTIVE SAR TACTICAL
MAPPER**



John B. Hampshire II

MAY 1997

Approved for public release; distribution is unlimited.

© 1997 Elsevier Science Ltd.
© 1997 John B. Hampshire II

This work contains copyrighted material. The United States has for itself and others acting on its behalf an unlimited, paid-up, nonexclusive, irrevocable worldwide license. Any other form of use is subject to copyright restrictions.

**SENSORS DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7318**

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) May 1997		2. REPORT TYPE Conference Paper		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE sarMAPPER: A REAL-TIME INTERACTIVE SAR TACTICAL MAPPER				5a. CONTRACT NUMBER F33615-97-1-1017	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 62301E	
6. AUTHOR(S) John B. Hampshire II				5d. PROJECT NUMBER ARPA	
				5e. TASK NUMBER AA	
				5f. WORK UNIT NUMBER IN	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University Department of Electrical & Computer Engineering Pittsburgh, PA 15213-3890				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Sensors Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7318				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL/SNAR	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-SN-WP-TR-2003-100	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES © 1997 Elsevier Science Ltd. © 1997 John B. Hampshire II This work contains copyrighted material. The United States has for itself and others acting on its behalf an unlimited, paid-up, nonexclusive, irrevocable worldwide license. Any other form of use is subject to copyright restrictions. Conference paper published in the <i>DARPA97 Image Understanding Workshop, May 1997</i> . This version of the report is the best quality available.					
14. ABSTRACT sarMapper is near-real-time, interactive software for generating high-accuracy tactical ground cover maps from synthetic aperture radar (SAR) imagery using current-technology laptop computers. Such maps make it possible to automate the focus-of-attention mechanism that is the foundation of tactical image analysis, target detection, and target recognition. This document describes a proof-of-concept first instantiation of sarMapper and outlines further technical issues to be addressed in the 1997 sarMapper research effort.					
15. SUBJECT TERMS SAR, Mapper, near-real-time interactive software					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 26	19a. NAME OF RESPONSIBLE PERSON (Monitor) Jason Johnson 19b. TELEPHONE NUMBER (Include Area Code) (937) 255-5668 x4047
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			

sarMapper: A Real-Time, Interactive SAR Tactical Mapper

John B. Hampshire II

Institute for Complex Engineered Systems (ICES)

and

Department of Electrical & Computer Engineering

Carnegie Mellon University

Pittsburgh, PA 15213-3890

Email: hamps@ece.cmu.edu

<http://gussolomon.ius.cs.cmu.edu/hamps/IU/index.html>

Abstract

sarMapper is near-real-time, interactive software for generating high-accuracy tactical ground cover maps from synthetic aperture radar (SAR) imagery using current-technology laptop computers. Such maps make it possible to automate the focus-of-attention mechanism that is the foundation of tactical image analysis, target detection, and target recognition. This document describes a proof-of-concept first instantiation of sarMapper and outlines further technical issues to be addressed in the 1997 sarMapper research effort.

1 Introduction

sarMapper is computational software and a human-machine interface that learns how to generate ground cover maps from synthetic aperture radar (SAR) images, given minimal human supervision. It is a real-time, interactive

learning tool that the military (as well as scientists) can teach and subsequently use to make maps for tactical (or scientific) analysis of ground cover. sarMapper runs on current-technology laptop computers, generating detailed mega-pixel maps in one to three minutes, depending on the level of map detail specified.

1.1 Background

Current-generation military SAR platforms produce vast quantities of imagery. Ground truth for this imagery — areas in which the type and quantity of ground cover is known in detail — is both scarce and expensive to obtain. Moreover, temporal changes in ground cover mean that ground “truth” is ephemeral. Added to these facts are tactical imperatives of military operations that rely on timely, accurate maps:

- military image analysts are drowning in a sea of data for lack of a real-time ability to process the data into usable information.
- military commanders need accurate tactical maps *now*, not two hours from now.
- they need them in the field, where their troops can use them; the troops need to be able to update their maps rapidly in order to reflect the changing tactical situation in real-time.
- they must be able to do this with minimal effort and training and little or no prior information regarding the area being imaged by the SAR reconnaissance platform.

sarMapper addresses these issues by learning to

This newly-funded research is sponsored by the Defense Advanced Projects Research Agency under grant F33615-91-1-1017, monitored by the United States Air Force Wright Laboratory ATR Development Branch, Wright Patterson AFB, Dayton, OH. The views and conclusions contained in this document are the author's and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Project Agency, Wright Laboratory, the U.S. Air Force, or the U.S. government.

generate detailed ground cover maps in near-real-time with minimal human intervention, using a current-technology laptop computer. The sarMapper described herein is a proof-of-concept. The proof-of-concept was designed to show that detailed ground cover maps can be generated using a human-guided "semi-supervised" parametric learning procedure without prior knowledge of ground cover. This paper describes the learning procedure, validates the concept, and describes the underlying technical and computational framework that make its real-time, interactive implementation possible.

1.2 sarMapper Overview

The sarMapper project seeks to provide the military with two fundamental capabilities:

- Fast, automated, high-accuracy, tactical map generation.
- Tactical focus of attention.

sarMapper itself is to have the following characteristics:

- Real-time learning & map generation
- Without prior ground-truth
- On a laptop (with a CD-ROM or large external disk)
- Computational efficiency (generate a high-accuracy mega-pixel map in one to three minutes on a laptop)
- High resolution / accuracy ground cover assessment
- Interactive graphical user interface (GUI): human aids computer in initial learning phase; after learning, computer maps autonomously
- Human oversight
 - Assess focus-of-attention warnings
 - During sarMapper's learning phases
- Multiple
 - Wavelength (P-, L-, C-, X-Band, etc.)
 - Polarization (single and fully polarimetric)
 - Spatial resolution
 - Data Sources (e.g.,...)
 - MSTAR public data (airborne, X-band)
 - LL ADTS SAR (airborne, X-band)
 - JPL AirSAR (airborne, P, L, and C-band)
 - NASA SIR-C/X-SAR (spaceborne L, C, and X-band)

- Focus-of-Attention (FOA)

- Detect & locate man-made ground cover
- Warn human according to prior tasking
- Usable with ~1 hour training

Computational efficiency forms the core of sarMapper, allowing it to generate mega-pixel maps on a current-generation laptop in one to three minutes. Ground cover types are learned using low-complexity parametric models of RF backscatter: learning takes the form of efficient model parameter estimation, which allows ground cover types to be characterized in terms of their backscatter signature — the learning and subsequent mapping take place in near real-time, owing to the efficient, low-complexity algorithms employed. In comparison, standard maximum-likelihood map generation algorithms generate maps on a time scale of hours.

sarMapper uses *semi-supervised* learning, which obviates the need for prior ground truth. The principle behind this supervised learning procedure is straightforward: humans can discern different ground cover types in a SAR image by the differences in their appearance in the image. Different ground cover types in a single-polarization image will appear to have different shades and/or textures of gray; in a false-color composite of multiple polarizations, different ground cover types will appear in different colors. Consequently, a human can identify regions of different ground cover in an image and label these regions without knowing what the different ground cover types are — using pseudonyms for the unknown ground cover classes. These pseudo-classes of ground cover can be learned, and their backscatter signatures can then be used to generate a high-resolution pseudo-map over a wide-area in the vicinity of the image used for learning. Many of the unknown ground cover types can be inferred by an image analyst from context, site-invariant backscatter signatures, historical imagery, or focussed follow-on surveys conducted using the pseudo-map to target specific survey sites. The critical characteristic of semi-supervised learning is that it can generate a useful map in real-time without prior knowledge of the area; missing details can be filled in as they are obtained, without having to re-learn or re-map the area.

Since man-made ground cover tends to backscatter little RF energy (as in the case of

obliquely illuminated metal or concrete surfaces) or substantial RF energy (as in the case of trihedral reflectors common to military vehicles), semi-supervised learning is compatible with sarMapper's focus-of-attention (FOA) mission. Very large areas (hundreds to thousands of square kilometers) can be mapped and surveyed for small areas of potential tactical interest by a single human and sarMapper team. Survey time for a thousand square kilometer area (10-meter map resolution) requires between ten minutes and one-hour on a single laptop computer, depending on the number of ground cover classes enumerated; survey time for a fifty square kilometer area requires between 15 seconds and three minutes under the same conditions.

1.3 Proof-of-Concept

The proof-of-concept objectives were to

- generate detailed ground cover maps from JPL AirSAR images *without* prior knowledge of ground cover.
- validate Tobit parametric models of radio frequency (RF) backscatter envelope at P, L, and C bands.
- show that these Tobit models can be learned with human supervision, even though the human has no prior knowledge of the ground cover at the site being mapped.
- show that Rayleigh-compressed SAR images (see section 2.2) and their associated Tobit-generated maps (section 2.5) can be used to target a small, well-defined, set of locales in a site (or set of sites) where ground truth should be surveyed.
- show that the Tobit-generated map can be converted to a true ground cover map with a few key strokes, once ground truth has been established at the selected survey sites.
- validate a number of image processing and computational advantages derived from the use of Tobit parametric models for ground cover.
- point the way to a highly automated focus-of-attention version of sarMapper.

2 Technical Summary

Figure 1 illustrates the sarMapper graphical user interface (GUI). In this paper, sarMapper takes SAR images that have been decompressed from the JPL AirSAR compression format described in

[DuBois-87,AirSAR-90], but sarMapper is designed to work with arbitrary SAR platforms. The main control panel has controls for specifying the receive and transmit polarizations to be synthesized when the complex floating-point SAR data is compressed into a real 8-bit integer image. Common polarizations (e.g., HH, HV, VV, Left Circular, and Right Circular) are synthesized (when fully polarimetric data are available) by efficient code, tailored to the specific polarization. Arbitrary polarizations are synthesized via a general compression-and-synthesis algorithm, which — in the case of JPL AirSAR — is a modified ANSI-C version of FORTRAN code supplied by Eric Rignot and Pascale DuBois of JPL, section 334. Tools for labeling SAR images are above the polarization controls; these are used in the semi-supervised learning procedure described in section 2.3.

A browser (upper right corner of figure 1) shows all the SAR data files currently loaded; when a file name is highlighted, the contents of its header are displayed in an information window, which allows string searches (so the user can quickly find information of interest). Multiple SAR images can be generated from the compressed SAR data files and displayed simultaneously. Figure 1 contains two false color images derived from C- and P-band AirSAR data over NASA's Raco, Michigan super site. Each of these images overlays HH (red), HV (green), and VV (blue) polarized images to form the color composite (see the author's IU web site for a color version of this figure). The lower-right window contains locale-specific histograms for each of these three polarizations in the P-band image: they and their corresponding parametric models, listed in the lower portion of the window, are related to the learning and map generating functions of sarMapper described in sections 2.1, 2.4, and 2.5.

2.1 Parametric RF Backscatter Models

sarMapper uses SAR amplitude (i.e., RF envelope) images to generate ground cover maps. Envelope statistics are used to derive parametric models of radio frequency (RF) backscatter; when the SAR is fully-polarimetric, three parametric models are derived for each ground cover class specified for a site: one model for each of the three polarizations (HH, HV, and VV) that comprise a color-composite image. One parametric model is generated when the SAR is

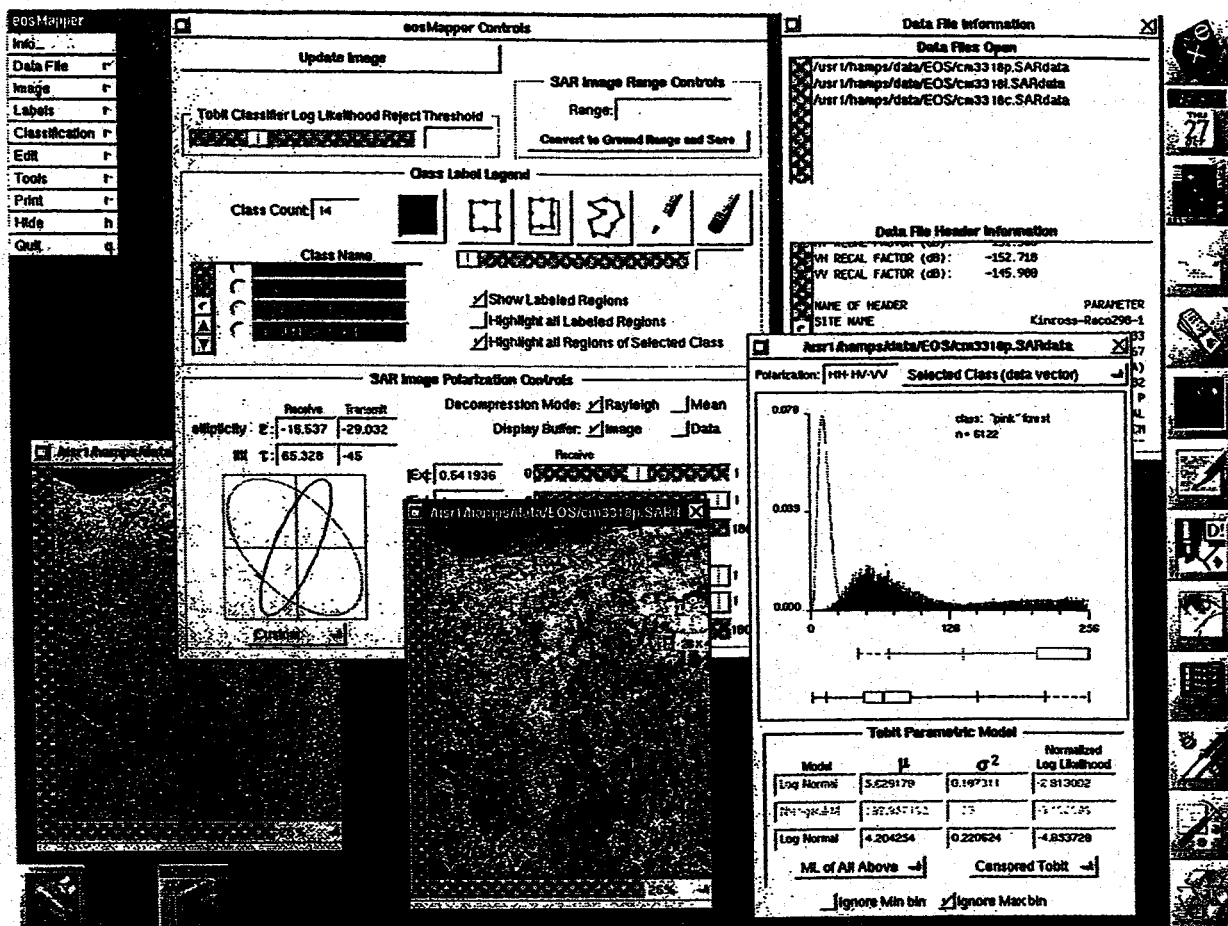


Figure 1: The sarMapper graphical user interface (GUI). The main controls allow the human operator to specify the polarization of the SAR image when the SAR is fully polarimetric, the number and type of ground cover classes, image labeling modes, etc. Multiple images can be loaded and displayed simultaneously (HH-HV-VV composites are shown for C- and P-band images of NASA's Raco Michigan super site imaged by the JPL AirSAR platform). Histograms of the backscattered SAR radio frequency (RF) envelope can be generated from the entire image, a single region, or all the regions belonging to a user-specified ground cover class: These histograms are used to generate Tobit maximum-likelihood parametric models for each ground cover class, which are in turn used to generate a ground cover map for the site and surrounding sites.

single-polarization. sarMapper models backscatter envelope (a real-valued scalar random variable); it does not model backscatter as a complex (in-phase and quadrature) random variable, as is common in the SAR literature (e.g., [Kong-87]). The reason for this is complexity. Parametric models are derived from one-dimensional histograms of RF backscatter when the envelope is used¹; when complex backscatter is used, the histograms are two dimensional. Assuming eight-bit precision, a (one-dimensional) envelope histogram has 256

bins, whereas a (two-dimensional) complex backscatter histogram has $256^2 = 65,536$ bins. Consequently, the training sample size necessary to generate a relatively smooth complex histogram — from which a low-bias, low-variance parametric model can be inferred — is, roughly speaking, the square of the size necessary to generate a relatively smooth envelope histogram and associated parametric model. Moreover, the computational complexities of evaluating log-likelihood functions — a necessary step in estimating parametric models — are proportional to the dimensionality of the histogram used: in short, using RF envelope, sarMapper avoids Bellman's curse of

¹See section 2.6 for the rationale behind the use of histogrammed versus full-precision amplitude statistics in the parametric modeling procedure.

dimensionality [Bellman-61,Duda-73]. The computational and sample complexities of the parametric modeling procedure are orders of magnitude lower than they are for the complex backscatter paradigm (see section 2.6), so the resulting parametric classifier is significantly more efficient in its data and computational requirements.¹

sarMapper uses Rayleigh, Rician (aka Rice-Nakagami), Nakagami-M, and LogNormal parametric models of the backscattered RF envelope. The Rayleigh, Rician, and Nakagami-M models are progressively general ones governing the statistics of the backscattered RF envelope when the complex backscatter is Normally distributed: a concise and well-formulated derivation of these models is given in [Beckmann-67]. The LogNormal parametric model spans a wide range of disciplines from econometrics to RF propagation; see [Strohbehn-75] for a motivation and description relevant to SAR backscatter. The probability density functions (pdfs) of these models follow, along with references: exhaustive details of all of them can be found in the appendices of [Hampshire-88]. In all the following pdfs α denotes the backscatter envelope random variable (rv), and A denotes the domain of α (i.e., $\alpha \in A$); $E_A[\alpha]$ denotes the expected value of α over its domain; and $f_\alpha(a)$ denotes the pdf of the rv α evaluated at $\alpha = a$.

Rayleigh pdf: α is distributed according to [Strutt-29,Davenport-58,Beckmann-67,Ishimaru-78]

$$f_\alpha(a) = \frac{a}{\sigma^2} \exp\left(-\frac{a^2}{2\sigma^2}\right), \quad (1)$$

where $2\sigma^2 = E_A[\alpha^2]$. Note that the Rayleigh pdf's mean is given by

$$E_A[\alpha] = \sqrt{\sigma^2 \cdot \frac{\pi}{2}} \quad (2)$$

¹Here I use the term "efficient" in both the classical Cramér-Rao context [Rao-45,Cramér-46] and the related pattern recognition context I have defined elsewhere [Hampshire-93a,Hampshire-93b].

Rician pdf: α is distributed according to [Rice-44,Nakagami-60,Beckmann-67]

$$(3) \quad f_\alpha(a) = \frac{a}{\sigma^2} \exp\left(-\frac{a^2 + a_0^2}{2\sigma^2}\right) I_0\left[\frac{a \cdot a_0}{\sigma^2}\right],$$

where a_0 is the value of the deterministic component of α , σ^2 is a variance-like parameter, and $I_0[\cdot]$ denotes the modified Bessel function of order zero [Abramowitz-70].

Nakagami-M pdf: α is distributed according to [Nakagami-60,Beckmann-67]

$$(4) \quad f_\alpha(a) = \frac{2m^m \cdot a^{(2m-1)}}{\Gamma(m) \cdot \Omega^m} \exp\left(-\frac{ma^2}{\Omega}\right),$$

where $\Omega = E_A[\alpha^2]$, m is an inverse normalized variance parameter, and $\Gamma[\cdot]$ denotes the gamma function [Abramowitz-70].

LogNormal pdf: α is distributed according to [Aitchison-66,Strohbehn-75]

$$(5) \quad f_\alpha(a) = \frac{1}{\sqrt{2\pi} a \sigma_{\ln(\alpha)}} \exp\left[\frac{(\ln(a) - \gamma)^2}{2\sigma_{\ln(\alpha)}^2}\right]$$

where $\gamma = E[\ln(\alpha)]$ (i.e., γ is the mean log-envelope) and $\sigma_{\ln(\alpha)}^2 = E_A[\ln(\alpha)^2] - \gamma^2$

(i.e., $\sigma_{\ln(\alpha)}^2$ is the variance of the log-envelope).

2.2 Rayleigh Image Compression & Equalization

SAR signals have dynamic range on the order of 55 dB², but an eight bit number encodes at most about 53 dB of dynamic range. Consequently, when a SAR image is displayed on a typical computer monitor, many of the pixels are

²Personal communication regarding the dynamic range of JPL AirSAR from J. Van Zyl, November, 1994.

saturated, taking on the minimal value of 0 or the maximal value of 255.

RF envelope (amplitude) displays of SAR imagery are formed by converting the complex backscatter amplitude (in-phase and quadrature components are generally represented with 32-bit floating-point representation or some derivative of this format) into a real magnitude represented with an 8-bit integer. If this compression is done with too much gain the resulting image is very bright; if it is done with too little gain, the resulting image is very dark. In either case, image detail is lost (either due to high-end or low-end saturation). The key to good image contrast and brightness is a compression scheme that minimizes the amount of both high- and low-end saturation while preserving as much of the dynamic range contained in the complex floating-point representation. From an information-theoretic perspective, the compression scheme that minimizes saturation and maximizes dynamic range minimizes the information loss between the complex floating-point and real integer image representations.

Typically, SAR images are compressed so that the mean pixel value is some target value in the vicinity of 128 (half of 8-bit full scale). The mean pixel value is estimated by sub-sampling the image with some default gain factor. Then the full image is compressed using a new gain factor that results in a mean pixel value close to the specified target value. This works well if the pixels are Normally distributed, but they are not. Rather the pixels are generally distributed according to one of the pdfs described in the previous section, all of which are characterized by a fairly high level of skewness. The histogram in the lower left image of figure 2 illustrates. It is based on the HH-polarized image generated from a JPL AirSAR L-band image of the Flevoland super site in The Netherlands (top, left). The image was generated using the standard mean-based compression scheme just described. The whole image histogram is approximately LogNormally distributed, with a very long tail: three quarters of the pixels have values less than 120, but more than ten percent of all pixels have the maximum value of 255. This is because the mean-based compression algorithm tends to set the compression gain too high in its attempt to match the mean pixel value with the target value (again, owing to the skewness of the RF envelope

pdfs described above).

Rayleigh compression works somewhat differently. Instead of trying to match the mean pixel value with some target near half scale, it determines the compression gain factor so that the fraction of pixels with the maximum value of 255 approximates a pre-determined target value of 5%. I call this value the "target saturation fraction". Matching the actual fraction of saturated pixels to the target fraction ensures that the image is neither too bright nor too dark. The compression scale factor needed to realize the target saturation fraction is computed by sub-sampling the image, fitting a Rayleigh pdf to the resulting histogram, and estimating the *cumulative* distribution function (cdf) for the data (under the Rayleigh assumption). The cdf can then be used to determine the compression scaling factor necessary to achieve the target saturation fraction. Here is how it works...

We have the sub-sampled image pixels, which we view as realizations of the RF envelope backscatter rv α . We compute the maximum-likelihood estimate of the Rayleigh pdf's parameter σ^2 in (1). Let's call this estimate $\sim\sigma^2$. Next we compute the Rayleigh cdf for the high-end saturation amplitude minus one ($\alpha = 254$) — which we denote by $\zeta_{\alpha}(254)$ — given $\sim\sigma^2$. This is given by

$$(6) \quad \zeta_{\alpha}(254) = 1 - \exp\left(-\frac{(255)^2}{2\sim\sigma^2}\right).$$

We can then compute the compression scaling factor δ that we should use so that the cdf of the re-scaled backscatter rv ($\delta\alpha$), evaluated at the saturation amplitude minus one (254), is equal to one minus the target saturation fraction f_s :

$$(7) \quad \zeta_{\delta\alpha}(254) = 1 - \exp\left(-\frac{(255)^2}{2\sim\sigma_{\delta}^2}\right) \equiv 1 - f_s.$$

Solving (7) for the Rayleigh parameter $\sim\sigma_{\delta}^2$, which corresponds to the target saturation fraction f_s , we obtain

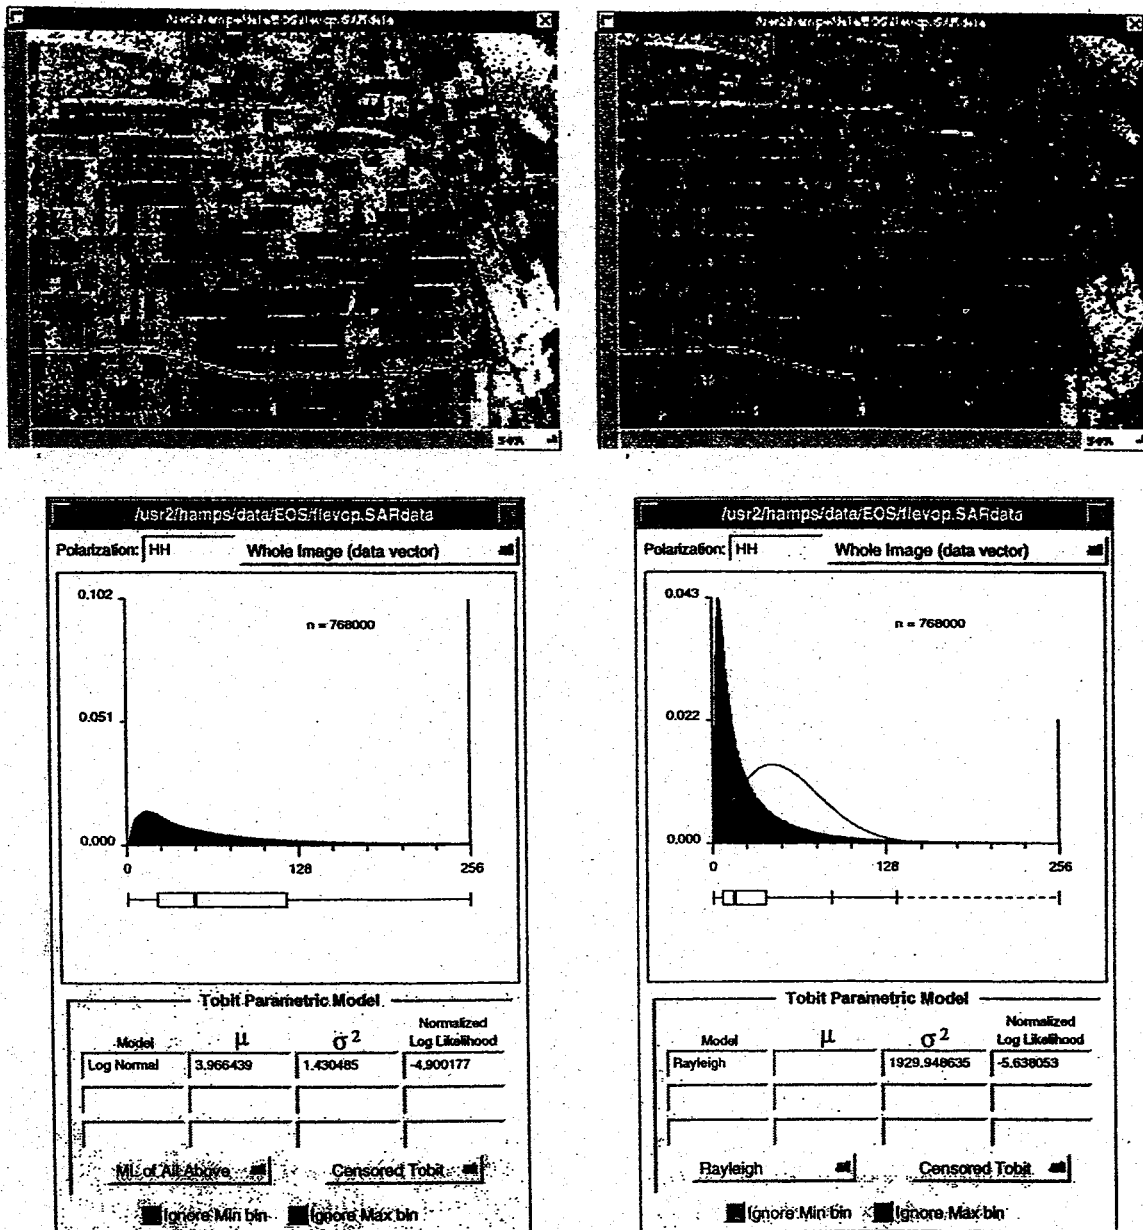


Figure 2: Two HH-polarized images generated from a JPL AirSAR L-band image of the Flevoland super site in The Netherlands. The left-hand image was generated using mean-based compression: the whole-image histogram below the image indicates that about ten percent of the image pixels are saturated (i.e., they have the maximum value of 255). The large number of light regions in the image confirms this. The right-hand image was generated using Rayleigh-based compression in which the image saturation was limited to Order[5%]: the whole-image histogram below indicates that only about two percent of the image pixels are saturated. The resulting monochromatic HH image is darker and shows less detail than its mean-compressed counterpart in this grayscale version, but the associated HH-HV-VV color composite image has better contrast and color balance. More importantly, the statistically consistent balance between image contrast and brightness afforded by Rayleigh compression equates to consistent dynamic range in the 8-bit data vector used for parametric learning and classification. This, in turn, leads to more consistent parametric models for ground cover classes and more accurate maps.

$$(8) \quad -\sigma_{\delta}^2 = -\frac{(255)^2}{2\ln(f_s)}.$$

Given (2), we find that the scaling factor that will generate f_s is given by

$$(9) \quad \delta = \frac{E_A[\delta \cdot \alpha]}{E_A[\alpha]} = \frac{-\sigma_{\delta}}{-\sigma} = \sqrt{\frac{-(255)^2}{2\ln(f_s) \cdot -\sigma^2}}; 0 < f_s < 1$$

The right-hand side of figure 2 shows the effect of Rayleigh compression. The scaling factor δ is computed for a target saturation fraction $f_s = .05$ and the image is compressed. After compression, the whole-image histogram is generated (bottom, right) and we find that the actual saturation fraction is about 2%. The difference between the target value of 5% and the actual value of 2% follows from the actual distribution of the whole-image data. In computing the scaling factor, we *assumed* that the whole image is Rayleigh distributed, when *in fact* it is approximately LogNormally distributed. Because the Rayleigh pdf isn't a proper model of the data, our computations are biased, and our actual saturation fraction is smaller than the 5% target.

This requires an explanation of our rationale for using only the Rayleigh pdf to compute the scaling factor instead of considering other models: *the Rayleigh cdf can be computed in closed form*. Consequently, the scaling factor δ can also be computed in closed form (i.e., very efficiently). What we lose in an imprecise estimate of the scaling factor necessary to achieve the target saturation fraction, we make up for in speed (the whole scaling factor estimation procedure for a mega-pixel SAR image takes a fraction of a second). The Rayleigh pdf is generally sufficient to get us within plus or minus 3% of our target saturation fraction of 5%.

Rayleigh compression generates images with *statistically consistent* dynamic range and saturation; mean-based compression generates images with dynamic range and saturation that can vary significantly with the skewness of the whole-image histogram. This statistical consistency contributes to statistically better ground cover maps generated from the Rayleigh compressed images.

One final note on the Rayleigh compression algorithm... The HH-polarized image is sampled to compute the scaling factor because it is consistently the one with the highest dynamic range. It is sub-sampled by sampling *all* the pixels in every 17th line in the image, beginning with the 23rd image line — these numbers are intentionally prime numbers, so the chance of computing a bad estimate of the whole-image statistics due to some periodic artifact is reduced.

2.3 Semi-Supervised Learning

sarMapper was conceived with two basic assumptions:

- There is usually little or no ground truth associated with a SAR site to be mapped.
- SAR backscatter statistics correlate strongly with the physical structure of the illuminated ground cover, and parametric models of the backscattered RF envelope are well understood (section 2.1).

The lack of ground truth implies an unsupervised learning procedure, but the strong prior knowledge we have about the statistics of radar backscatter suggests a parametric modeling paradigm, which involves supervised learning. I concluded that it would be possible to design a "semi-supervised" learning procedure that would use parametric models without prior knowledge of ground cover. What I will describe in this section is the first instantiation of this semi-supervised learning, which involves a human supervisor. The critical reader will note that human-supervised learning is a risky business because the supervision is subjective (we are non-deterministic creatures) and it varies from human to human (we are all different). Statistically, therefore, the level of human involvement in this first instantiation of semi-supervised learning is a weakness, but it conveys two advantages:

- It exploits the strong inferential capabilities of the human supervisor.
- It is fast and highly interactive. The user can interact with the machine and see the results of this interaction in real-time. For example, he can make a change to the ground cover regions specified for the machine's learning phase, have the machine re-learn the regions,

and generate a new map — all in about two minutes.

How does the human user identify these homogenous regions? sarMapper generates a color composite image of the site by overlaying an HH image in red, an HV image in green, and a VV image in blue in the case of fully-polarimetric radar (or a single grayscale image of the HH polarization in the case of single-polarization radar); figure 3 shows such an image, derived from a JPL AirSAR P-band image of the Landes maritime pine forest in France. The user then identifies a number of regions in the image with the same basic color (or shade of gray), outlines them or colors them in, and groups them into one so-called “pseudo-class”. This process is repeated for each pseudo-class identified in the composite image by the human supervisor: each pseudo-class is assigned a map color and a pseudonym. sarMapper’s main control panel, shown in figure 1, contains the region labeling controls. Regions representing a given ground cover pseudo-class can be outlined using one of three modes (square region, rectangular region, or polygon) or they can be “penciled in” — a fourth labeling mode that identifies small, oddly-shaped regions (e.g., winding rivers) by coloring in a local binary bitmap. The number of ground cover classes, their names, and their map colors are user-specified to the right of the labeling controls in figure 1.

Once the training regions for each pseudo-class have been identified by the human supervisor, sarMapper characterizes each ground cover pseudo-class by computing the maximum-likelihood parameters for all four possible backscatter parametric models (Rayleigh, Rician, Nakagami-M, and LogNormal) — it does this for each of the three polarizations (HH, HV, and VV). The parametric model with the greatest log-likelihood (summed across the three polarizations) is then chosen as the maximum-likelihood (ML) model for the ground cover pseudo-class. Summing the log-likelihoods implicitly assumes that the HH, HV, and VV polarized images are independent, which they are not; the procedure is taking a statistical liberty for the sake of computational simplicity. The resulting ML model characterizes its associated ground cover pseudo-class and is used — along with all the other ML models — to produce a map (section 2.5). Figure 3 shows how the color

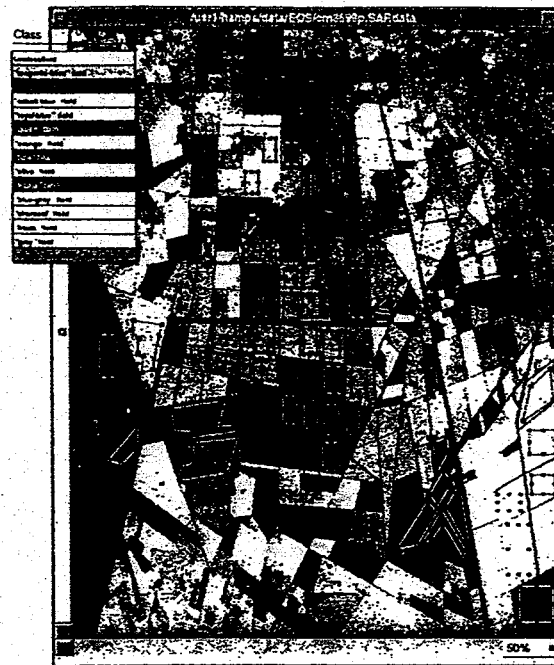


Figure 3: A sarMapper-generated HH-HV-VV false-color composite image of the Landes forest in France. The composite (shown here in grayscale — see the author’s IU website for a color version) was generated from a JPL AirSAR P-band image. A human labels distinctly colored regions in the case of a fully-polarimetric radar (or shaded ones in the case of a single-polarization radar) denoting distinct backscattering properties in the image. These regions are shown above with “hooks” used to move/resize them. The regions are given color codings and pseudonyms (since the actual type of ground cover is generally *not* known *a priori*). The legend in the upper left corner of the figure shows these pseudonyms and color codings. sarMapper then learns a Tobit maximum-likelihood parametric model for each ground cover class and generates a color-coded ground cover map of the entire site and its surrounding areas.

composite image looks when a number of pseudo-classes (note the names) have been identified and outlined. At this point, sarMapper is ready to generate Tobit ML parameter estimates for all the candidate pdfs, in order to determine the ML model for each pseudo-class.

In the next two sections I will describe how Tobit parametric estimation works, how Tobit models

are used to generate ground cover pseudo-class maps, and how these pseudo-maps are converted into true ground cover maps.

2.4 Tobit Estimation of Ground Cover Class-Conditional Densities with RF Envelope Backscatter Models

In section 2.1 I described the parametric models that sarMapper uses to characterize the backscattered RF envelope and I hinted why sarMapper uses eight-bit data representations and real backscatter (amplitude) histograms, rather than 32-bit floating point data representations of complex backscatter (I explain this fully in section 2.6). In section 2.2 I explained that Rayleigh compression is used in order to control the amount of saturation in the SAR images synthesized during compression of the JPL AirSAR data. Because the backscatter contains at least 2 dB more dynamic range than can be encoded with an eight bit number, there is *always* some image saturation, which raises the following question: how does sarMapper develop good parametric models of regions that generate heavily saturated RF backscatter statistics? The answer: by using a Tobit (i.e., Tobin probit) maximum-likelihood parameter estimation procedure.

A detailed explanation of Tobit estimation is well beyond the scope of this paper, so I will discuss it in very general terms. Tobit estimators were first developed in the 1940's by Hald [Hald-49], but the work went un-recognized. When it was developed independently and published almost a decade later by Tobin [Tobin-58], it was called the "Tobin-probit" model, a name that was later shortened to "Tobit model". Amemiya has written extensively on Tobit models in the econometrics context; his most general description can be found in [ch. 10, Amemiya-85]. I describe Tobit models in the RF backscatter context in [Hampshire-88, Hampshire-92]. Given a random variable α with a known pdf, the corresponding Tobit model is the pdf for the random variable α' , which α becomes when it is measured by a saturating device. In its simplest form, a high-end saturating device produces α' by "clipping" α for all values above the saturation threshold α_τ :

$$\alpha' = \begin{cases} \alpha, & \alpha < \alpha_\tau \\ \alpha_\tau, & \text{otherwise} \end{cases} \quad (10)$$

Audiophiles will recognize the clipped output of an over-driven amplifier as a saturated random variable. The audio mathematics are more complicated because the transition from linear to non-linear amplifier operation is gradual, *not* discontinuous, as it is in (10).

When computer images are generated from SAR data, the backscattered RF envelope α is converted to a saturated, eight-bit-quantized envelope α' . Equation (10) is a reasonable description of the saturated RF envelope prior to eight-bit quantization. Rayleigh compression limits the amount of saturation in the whole image, but it does not limit the amount of saturation in localized regions of the image. As a result, some regions within an entire SAR scene are heavily saturated. Figure 4 shows the histogram for a heavily saturated region in a JPL AirSAR P-band image of the Raco, Michigan super site (HH-polarization). The left view shows the full histogram, and the right view shows the un-saturated part of the histogram. Thirty four percent of the region's pixels are saturated, so the histogram is dominated by the statistics of its most significant bin ($Q(\alpha') = 255$, where $Q(\cdot)$ denotes the 8-bit quantization operator). This is clear in the left-hand view. If we ignore the most significant bin, we see that the un-saturated part of the histogram still has a recognizable shape. When a LogNormal Tobit estimator is applied to the full, saturated histogram, the resulting parametric model fits the data well (a plot of the model is superimposed on the histogram in the right-hand view). Mathematically, the Tobit estimator combines what can be inferred from both the un-saturated and saturated parts of the histogram to estimate the parametric model of the *un*-saturated rv α from its quantized, saturated counterpart $Q(\alpha')$. In fact, Tobit estimators for the parametric models described in section 2.2 are *efficient* maximum-likelihood estimators (in the Cramér-Rao sense [Cramér-46, Rao-45]) [Hampshire-92]. They yield good models even when the amount of saturation is substantial (as it is in figure 4). Since sarMapper learns Tobit parametric models of the 8-bit quantized RF backscatter envelope to characterize each ground cover class, it is able to




Figure 4: Two views of one histogram, which characterizes the RF envelope (JPL AirSAR, P-band, HH-polarized) backscattered from a region in NASA's Raco, Michigan super site. The left-hand view shows the full histogram: about 34% of the HH image pixels in the region are saturated. The right hand view shows the un-saturated part of the histogram (i.e., the histogram for bins 0 – 254). About 66% of the pixels fall in this range. The Tobit ML estimation procedure found that the histogram was best fit by the LogNormal pdf using the parameters listed below the display. Note that the model (superimposed on the histogram) fits the histogram well, despite the heavy amount of saturation. This illustrates two important facts: an 8-bit histogram of the RF backscatter envelope remains a sufficient statistic for inducing the parametric model of the backscatter despite substantial saturation; the Tobit estimator remains efficient (in the Cramér-Rao sense) under substantial saturation. A standard maximum-likelihood parameter estimation procedure for the LogNormal pdf would yield a very poor fit to this saturated data (see [Hampshire-92]).

generate good maps from eight-bit amplitude images in near real-time.

The histograms and box plots [Tukey-77] for each ground cover pseudo-class are generated by compiling statistics on the values of all the pixels belonging to that class: this is done for the three (fully-polarimetric radar) or single (single-polarization radar) 8-bit (i.e., 256-value) plane(s) in the sarMapper data vector. For fully-polarimetric radar these correspond to the Rayleigh-compressed HH (red), HV (green), and

VV (blue) polarized data vectors for the image; for single-polarization the Rayleigh-compressed HH image is displayed in a single grayscale image plane. The ML Tobit model is then computed for each of the pseudo-classes as described in the previous section. The Tobit estimation procedures described in [Hampshire-88,Hampshire-92] govern the computations by which each model's parameters and log-likelihood are estimated.

The training data and computational

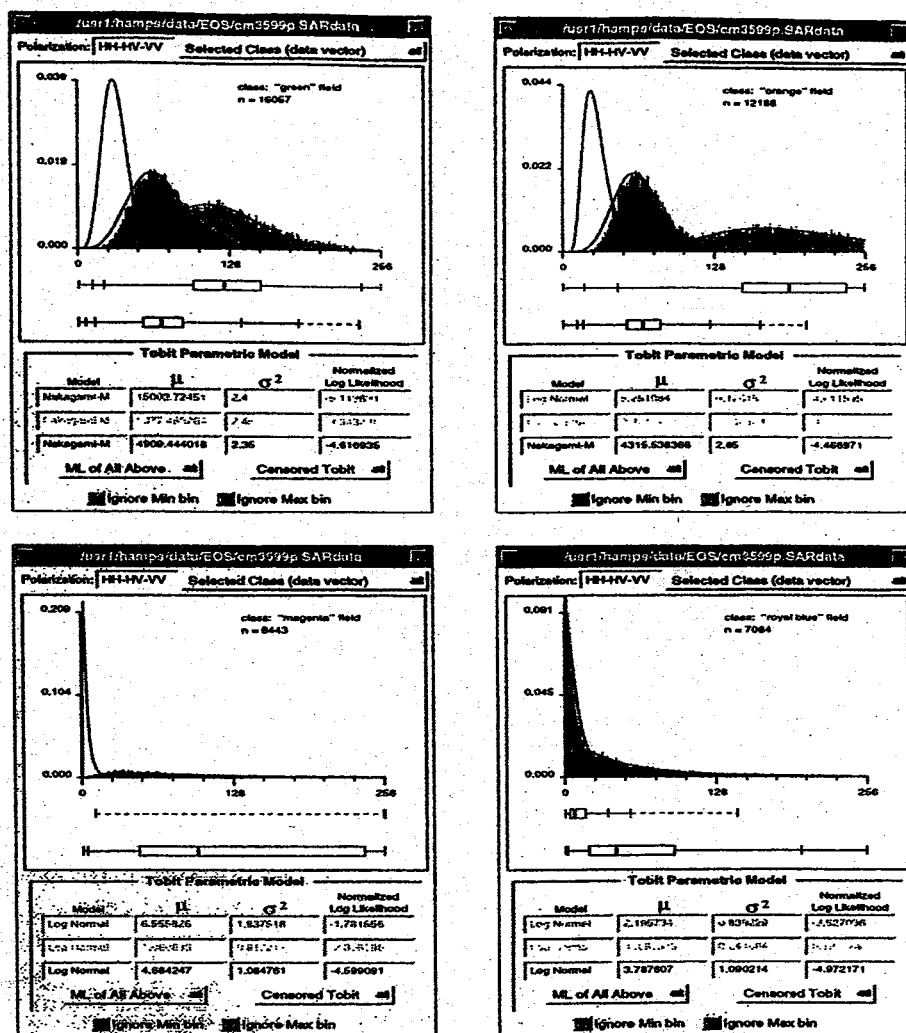


Figure 5: Histograms, box plots [Tukey-77], and the resulting Tobit maximum-likelihood (ML) parametric models for four ground cover pseudo-classes in the Landes maritime pine forest P-band image (see figure 3). The box plots and list of Tobit ML models are arranged as follows for each pseudo-class: HH (top), HV (middle), VV (bottom). The histogram polarizations are not obvious in a grayscale version of this figure, but can be inferred from their corresponding box plots. In general the HH histogram is the right-most (i.e., the one with the largest mean/median/mode), the VV is the central one, and the HV is the left-most (i.e., the one with the smallest mean/median/mode) for each pseudo-class. The Tobit ML parametric model for each polarization is superimposed as a solid line on its associated histogram. Nakagami-M and LogNormal models are the common (although not exclusive) ML choice since they represent more general RF scattering paradigms. Note that the models typically fit the histograms well. The Tobit parametric estimators are robust, despite the saturation that inevitably occurs when a radar signal is modeled with only 8-bit precision.

requirements of sarMapper are reduced by about four orders of magnitude when the low-complexity eight-bit/real data representation I have described is used instead of the 32-bit/complex representation commonly used. Tobit models allow sarMapper to exploit the efficiencies afforded by the low-complexity data representation without significant reductions in map accuracy or precision. Figures 5 illustrates

why this is so: it shows the maximum-likelihood Tobit parametric models for four ground cover pseudo-classes found in the JPL AirSAR P-band image of the Landes maritime pine forest (figure 3). In general, the Tobit models fit their class-conditional histograms well. When a pseudo-map is generated from these models and subsequently converted into a true ground cover map (using the procedure described in the next

section) figure 6 results. My initial comparison of this map with ground truth indicates an accuracy on the order of 80-90% for those areas in which ground truth is known, but the confidence bounds on this estimate are large since the amount of ground truth for the Landes site is limited. This result and the others like it suggest that sarMapper's use of the eight-bit/real representation and Tobit models is sufficient to generate high-accuracy maps in near real-time. Research plans for this calendar year include quantitative evaluation of this hypothesis.

2.5 The Tobit Parametric Classifier & Map Generation

The Tobit parametric classifier is nothing more than a set of Tobit models — one for each ground cover pseudo-class identified by the human user. The classifier generates a pseudo-map for the SAR site by computing the class-conditional log likelihoods for all the pixels in each $n \times n$ pixel region in the site (these regions are contiguous, and the user specifies the map resolution n): the class of the Tobit model with the largest log-likelihood is chosen as the class for the region. This is standard maximum-likelihood parametric classification.

The resulting map is a pseudo-map in that the human user is assumed to have no prior information about the site's true ground cover. If ground truth exists (or can be inferred), the human user cross matches areas with known ground cover to the corresponding ground cover pseudo-classes. For cases in which the correspondence is one-to-one, the user can convert the pseudonym to the true ground cover name by simply re-typing the class name. For cases in which a pseudo-class represents more than one ground cover class, the user replaces the pseudonym with a list of all the true ground covers that the pseudo-class represents. Map colors can also be changed interactively.

In my experiments, I generated a pseudo map and saved the labeled pseudo-class regions. When I obtained ground truth for the site, I loaded in the pseudo-map labeled regions and changed the names/colors to reflect the truth. I then saved this ground truth data in a second label file and used it to re-generate the ground cover map — this time with the correct ground cover names and, perhaps, more meaningful colors. Representative

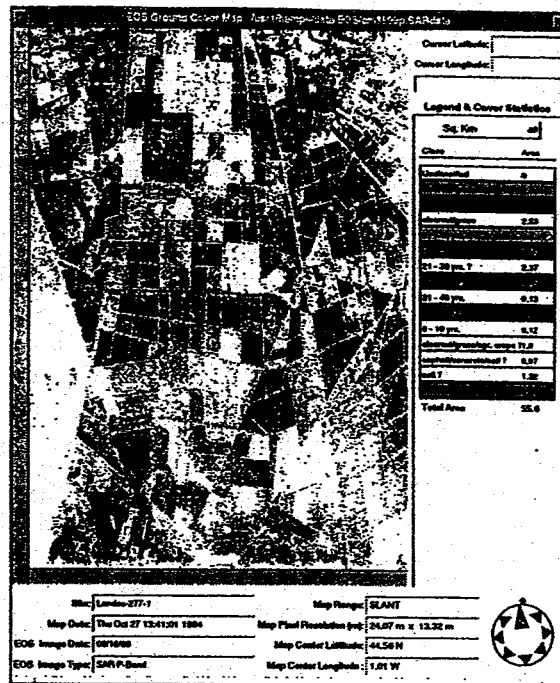


Figure 6: A ground cover map generated from the fully polarimetric image in figure 3. Fourteen different ground cover classes are distinguishable; of these seven represent the same species of tree at various ages, and four represent clearcut areas in various stages of regrowth. A full-color copy of this map can be found on the author's IU website (see references). Man-made ground cover is denoted by white and light gray in the lower/left part of this slant-range map.

maps are discussed further in section 3.

2.6 Efficiency Of The Tobit Maximum-Likelihood Classifier

Tobit maximum-likelihood mapping is both statistically efficient [Hampshire-92] and computationally efficient. Single-polarization JPL AirSAR images generally contain about 1.3 million pixels, so a 3-plane sarMapper color composite image (HH-HV-VV polarizations) contains about 4 million pixels. Using a 32-bit representation for each of the three polarizations, we would require almost 32 megabytes (Mb) of storage for each composite image. The requirement is 32 Mb rather than 16 because each data point would be a complex number. Let's assume that there are 10 ground cover pseudo-classes associated with an image. Using the 32-bit floating point representation, we would need

to evaluate 40 million log-likelihood functions — one for each polarization of each pixel, given each ground cover pseudo class — to generate a high-resolution map.

By using an 8-bit envelope (real) representation for each of the three image planes, we require only 4 Mb of storage per composite image. Moreover, if we have 10 ground cover pseudo-classes, we need only evaluate $3 \times 2^8 \times 10 = 7,680$ unique log-likelihoods — one for each polarization of each possible pixel value, given each ground cover pseudo-class. *sarMapper pre-computes* these log-likelihoods and writes them into a look-up table prior to generating the map: when it actually generates the map, it does table look-up to get the log-likelihoods (which is about four orders of magnitude faster than actually computing them). This is why *sarMapper* can generate a high-resolution map with 15 ground cover pseudo-classes in about one minute on a current-technology laptop computer.

Again, the Tobit parametric estimation procedure enables *sarMapper* to exploit the computational efficiency of the low-complexity representation with little or no reduction in map accuracy. In fact, map accuracy might even be *better* using the eight-bit/real representation instead of the 32-bit/complex one. This possibility exists owing to the Bellman's curse of dimensionality argument in section 2.1. Research plans for this calendar year include quantitative evaluation of this possibility.

3 Preliminary Results

Figures 8 and 9 are *sarMapper*-generated maps of NASA's Raco, Michigan supersite. They were generated from the JPL AirSAR C-band image in figure 7, using only the HH-polarization image (not shown) in order to approximate the process of generating maps from single-polarization X-band military imagery.

I spent about five minutes studying the image in figure 7, selecting regions of vegetation, water, and tarmac, which I inferred from the image without any ground truth. I generated the low-resolution map of figure 8 in approximately five seconds, and the high-resolution map of figure 9 in approximately 15 seconds. These two maps illustrate how *sarMapper* might be used to focus the attention of image analysts searching for




Figure 7: An HH-HV-VV composite C-band JPL AirSAR image of NASA's Raco, Michigan supersite (shown here in grayscale — see the author's IU website for a color version). Note the airfield in the lower right quadrant of the image.

man-made ground cover of potential tactical interest: figure 8 represents a first-stage focusing process that identifies ground cover that could be either water or tarmac. This very coarse map identifies a small number of regions in the map that warrant further analysis. Figure 9 is a higher-resolution of the site. Imagine that figure 8 is used as a focus-of-attention (FOA) mask overlayed on figure 9. Detailed mapping and analysis of the masked regions (figure 9) would reveal the area containing the airfield with a small number of additional "false tarmac" areas. A simple automatic target recognition (ATR) post-processing of this image would recognize the "false tarmac" regions as water, and the airfield as the single area of potential tactical interest.

Figure 10 shows another map of the Raco, Michigan super site, produced from a longer-wavelength JPL AirSAR P-band image. Knowing nothing about the site, I generated a pseudo-labeled HH-HV-VV color composite image and sent it to Leland Pierce (U.




Figure 8: A coarse(300 m) resolution focus-of-attention map of the Raco, Michigan site, generate from a C-band HH-polarized image. Darker regions on the map indicate areas of potential interest (water or tarmac; the latter is of tactical interest). High-resolution mapping and Automatic Target Recognition can be focussed on these areas.

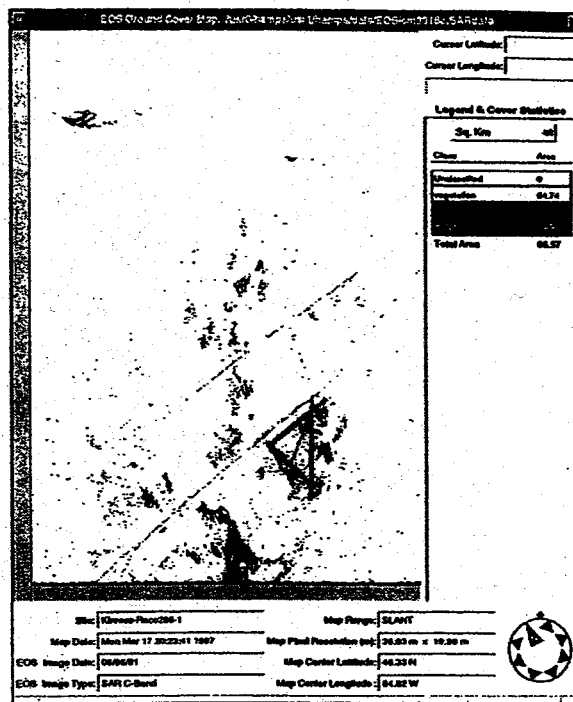


Figure 9: A fine(36 m) resolution map of the Raco, Michigan site, generate from the same C-band HH-polarized image used to generate the focus-of-attention map in figure 8. The darkest regions on the map indicate areas likely to be tarmac. The physical structure of the airfield, combined with its tarmac ground cover identify it as a legitimate target of interest.

Michigan), who graciously sent me ground truth corresponding to the regions I had sent him. I compared his ground truth to the map I had generated from the same labels and derived the true ground cover classes (shown) from the original pseudo-classes (see the author's IU web site for a color version of this map). The map is generally in good agreement with the ground truth, with some notable exceptions I discuss below. Five pseudo-classes corresponded to clearcut/grassland, two to non-scattering ground cover (water and man-made materials), two to hardwoods (aspen, birch, and northern hardwoods), and four to conifers. The P-band image does not discriminate non-scattering ground cover (water and man-made materials) from scattering ground cover well. This is due in part to the steep SAR incidence angle (small look angle) at the top of the image, which generates strong HH backscatter from all ground cover. Also, the long wavelength generates some backscatter from undulations in water surfaces.

The shorter wavelength in figure 7 generated substantially less backscatter from the water and man-made materials, so they were easier to distinguish from vegetation than they were in the P-band image. Water and asphalt/concrete should be more distinguishable in an X-band image. Nevertheless, figure 10 illustrates the utility of longer-wavelength radar for generating detailed wide-area tactical maps. A military commander planning an assault on the airfield in this image would know where to conduct paratrooper air-drops (in the fresh clear-cut areas to the north of the airfield), and would realize that an armored assault landed on the beach at the top of the map should *not* take a direct route towards the airfield. Such a route would require the armor column to penetrate a forest of old-growth hardwood trees (more than 20 meters tall). A less direct route eastward across the top of the map, and then south-southwest to the airfield would be better, since it would be through groves of aspen,

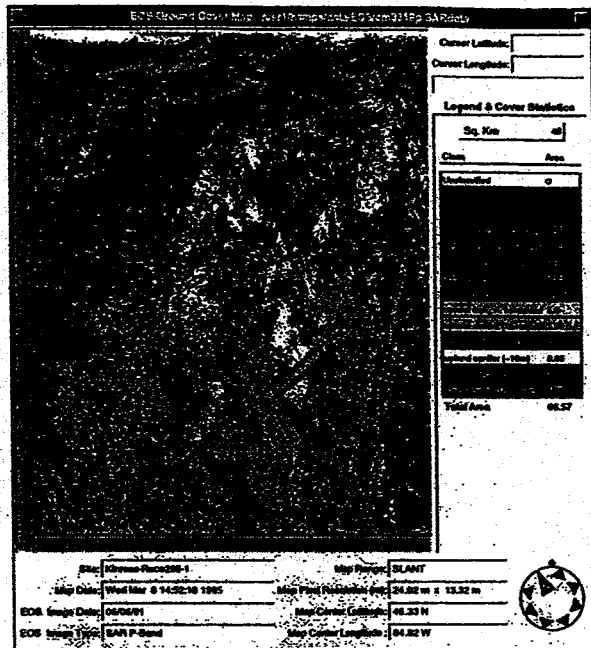


Figure 10: A detailed vegetation map of the Raco, Michigan site generated from a long-wavelength JPL AirSAR P-band image (see the author's IU web site for a color version of this map). The map indicates that sarMapper can distinguish subtle differences in ground cover, providing important tactical information to a commander who might, for example, be planning an assault on the airfield in the lower-right quadrant of the map (see text).

birch, and small conifers, which would be easier to penetrate.

4 Future Work

Pre-processing will be added to sarMapper so that it can map imagery from these three sensors (in addition to the JPL AirSAR platform):

- Lincoln Lab ADTS SAR (X-band)
- MSTAR SAR (Sandia/DOE X-band)
- NASA SIR-C/X-SAR (L, C, X-band)

Four research questions will be addressed this year. These questions — all of them technical in nature — touch on the speed, accuracy, and robustness of sarMapper's mapping and focus-of-attention capabilities:

- Can high-accuracy maps be generated from 8-bit X-band backscatter envelope data? Is the

accuracy of these maps significantly better or worse than the accuracy of maps generated from the same imagery using standard maximum-likelihood techniques on floating-point representations of *complex* RF backscatter?

- Is semi-supervised learning a statistically consistent paradigm?
- Can man-made ground cover be identified consistently in SAR imagery and pseudo-maps without ground truth?
- How can a robust focus-of-attention algorithm be derived for real-time laptop implementation?

4.1 Evaluation

Pursuant to research efforts to address them, the research questions listed in the previous section will be answered by objective evaluation of sarMapper using SAR data from a wide variety of sensors.

Speed: Average semi-supervised learning time will be assessed with human-computer timing trials. Map generation times will be tabulated for a corpus of evaluation images.

Accuracy: Map accuracy will be assessed using imagery for which ground truth is known or can be inferred. Accuracy will be quoted according to general methods of statistical inference/pattern recognition, with 95% confidence bounds and ground cover confusion matrices derived from test images (or test areas within an image) not used during semi-supervised learning.

FOA: sarMapper's focus-of-attention algorithm will be assessed according to general methods for evaluating detection algorithms; namely, receiver operator characteristic (ROC) curves will be generated and evaluated for FOA performed on test imagery/maps not used for semi-supervised learning.

Acknowledgements

Some funding for this proof-of-concept research was provided by NASA in 1994. Victoria Gor (JPL) translated my original Tobit estimation source code from FORTRAN into ANSI C. Dr. Cynthia Williams (Institute of Northern Forestry) has graciously volunteered some of her time to evaluate sarMapper's potential as a useful ecological mapping and analysis tool. Her insights have proven useful to the DARPA thrust

of this research.

Eric Rignot and Pascale DuBois of JPL section 334 provided me with FORTRAN source code and the technical details of JPL AirSAR data compression/decompression algorithms¹, which I converted to ANSI-C and subsequently modified. They and Jakob Van Zyl also provided me with a number of AirSAR images and technical papers they have written on the subjects of SAR and SAR remote sensing. Leland Pierce (University of Michigan) provided me with ground truth for the Raco, Michigan super site, as well as a number of helpful insights. Thuy LeToan (University Paul Sabatier, Toulouse, France) provided me with ground truth for the Landes maritime pine forest in France. Tony Freeman (JPL, section 334) provided me with AirSAR images and ground truth for the Flevoland site in Holland.

References

The author's DARPA IU website can be reached at

<http://gussolomon.ius.cs.cmu.edu/hamps/IU/index.html>

- [Abramowitz-70] M. Abramowitz and I. Stegun, editors. *Handbook of Mathematical Functions*, volume 55 of *U.S. Dept. of Commerce, National Bureau of Standards, Applied Mathematics Series*. U.S. Government Printing Office, Washington, D.C., 1970. Ninth printing.
- [AirSAR-90] *Airsar Bulletin*. Jet Propulsion Laboratory, Pasadena, CA, February 22 1990.
- [Aitcheson-66] J. Aitcheson and J. A. C. Brown. *The Lognormal Distribution*. Cambridge University Press, London, 1966.
- [Amemiya-85] T. Amemiya. *Advanced Econometrics*. Harvard University Press, Cambridge, MA, 1985.
- [Beckmann-67] P. Beckmann. *Probability in Communication Engineering*. Harcourt Brace and World, New York, 1967.
- [Bellman-61] R. E. Bellman. *Adaptive Control Processes*. Princeton University Press, Princeton, NJ, 1961.
- [Cramér-46] H. Cramér. *Mathematical Methods of Statistics*. Princeton University Press, Princeton, NJ, 1946.
- [Davenport-58] W. B. Davenport, Jr. and W. L. Root. *An Introduction to the Theory of Random Signals and Noise*. McGraw Hill, New York, NY, 1958. Re-printed by the IEEE press in 1987, ISBN 0-87942-235-1.
- [DuBois-87] P. DuBois et al. "Data volume reduction for imaging radar polarimeter" [sic]. In *Proceedings of IGARSS-87*, 1987. Abridge version in *NASA Tech Briefs*, Vol. 13, No.2, item 37.
- [Duda-73] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. John Wiley & Sons, New York, NY, 1973.
- [Hald-49] A. Hald. "Maximum Likelihood Estimation of the Parameters of a Normal Distribution which is Truncated at a Known Point." *Skandinavisk Aktuarietidskrift*, 32:119-134, 1949.
- [Hampshire-88] J. B. Hampshire II. *A Non-Rayleigh Model for Ultrasonic Backscatter in Myocardium*. Master's thesis, Theyer School of Engineering, Dartmouth College, August 1988.
- [Hampshire-92] J. B. Hampshire II and J. W. Strohbehn. "Tobit maximum-likelihood estimation for stochastic time series affected by receiver saturation." *IEEE Transactions on Information Theory*, IT-38(2):457-469, March 1992.
- [Hampshire-93a] J. B. Hampshire II. *A Differential Theory of Learning for Efficient Statistical Pattern Recognition*. PhD thesis, Carnegie Mellon University, September 1993.
- [Hampshire-93b] J. B. Hampshire II and B. V. K. Vijaya Kumar. "Differentially Generated Neural Network Classifiers are Efficient." In C. A. Kamm, G. M. Kuhn, B. Yoon, R. Chellappa, and S. Y. Kung, editors, *Neural Networks for Signal Processing III: Proceedings of the 1993 IEEE Workshop*, pages 151-160, New York, September 1993. The Institute of Electrical and Electronic Engineers, Inc.
- [Ishimaru-78] A. Ishimaru. *Wave Propagation and Scattering*, vol. 1. Academic Press, New York, NY, 1978.

¹These JPL AirSAR compression/decompression algorithms are unrelated to the algorithms described in section 2 (which relate to 8-bit displayable representations of SAR imagery); instead they relate to efficient storage of the Stokes matrix coefficients used as the fully-polarimetric AirSAR data representation.

- [Kong-87] J. A. Kong, A. A. Swartz, H. A. Yueh, L. M. Novak, and R. T. Shin. "Identification of Terrain Cover Using the Optimum Polarimetric Classifier." *Journal of Electromagnetic Waves and Applications*, 2(2):171-194, 1987.
- [Nakagami-60] M. Nakagami. "The m-distribution — a general formula of intensity distribution for rapid fading." In W. C. Hoffmann, editor, *Statistical Methods in Radio-Wave Propagation*, chapter 1, pages 3-36. Pergamon Press, New York, 1960. A summary of the author's work, contemporary with that of S. O. Rice, but relatively unknown in the U.S. until this publication.
- [Rao-45] C. R. Rao. "Information and accuracy attainable in the estimation of statistical parameters." *Bulletin of the Calcutta Mathematical Society*, 37:81-91, 1945.
- [Rice-44] S. O. Rice. "Mathematical Analysis of Random Noise." *Bell System Technical Journal*, 23:283-332, 1944. First of two parts.
- [Strohbehn-75] J. W. Strohbehn, T. Wang, and J. P. Speck. "On the Statistics of Line-of-Site Fluctuations of Optical Signals." *Radio Science*, 10:59-70, January 1975.
- [Strutt-29] J. W. Strutt. *The Theory of Sound*. MacMillan Co., London, 1929. The author is better known as Lord Rayleigh.
- [Tobin-58] J. Tobin. "Estimation of Relationships for Limited Dependent Variables." *Econometrica*, 41:24-36, 1958.
- [Tukey-77] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, Reading, MA, 1977.

sarMapper

Recent Enhancements


1. Ported sarMapper to a 166MHz Pentium laptop.
2. Extended capability to process SAR data from following sources.
 - ADTS (Lincoln Lab)
 - MSTAR (Sandia)
 - JPL AirSAR (NASA JPL: P, L, & C-band)
 - SIR-C (NASA JPL: L & C-band *work not yet complete*)
3. Added software to project slant range ground cover maps into ground range.

Summary of Initial sarMapper Results for ADTS/MSTAR data

1. Have run qualitative assessment of sarMapper on detecting MSTAR targets and shadow against background clutter (clutter from the target image chips). Initial results demonstrate ability to discriminate MSTAR targets and their associated shadows from surrounding ground cover.
2. Have run qualitative assessment of sarMapper with short wavelength HH-polarized SAR (compared with short, medium, and long wavelength HH-HV-VV polarized SAR). Results are encouraging on JPL AirSAR, ADTS, and MSTAR clutter data: we can still map a reasonable variety of ground cover with only HH-polarized X-Band data, e.g.,
 1. possible target
 2. shadow
 3. grass
 4. shrubs/trees
 5. pavement/dirt (not in all imagery)Longer wavelength data (e.g., JPL AirSAR P-band) allows greater discrimination of ground cover classes using only HH-polarization.
3. Embedded an MSTAR target in an ADTS scene (HH polarization only) and ran sarMapper over image to map it/detect possible targets as a demonstration of sarMapper's inherent focus-of-attention capability. Initial result looks good. T72 target was detected along with five other potential targets in the scene: the possible targets' total ground cover area was less than .5% of the total image area. Sensitivity of the FOA is easily modulated by varying sarMapper's resolution (which can be done in real-time with the push of a slider).


NOTE: These results are preliminary and qualitative only. Detailed quantitative analysis will follow this fall.

Examples of Initial sarMapper Results for ADTS/MSTAR data



The left-hand (or top) image above is that of a T-72 tank imaged by the MSTAR X-band (HH-polarization) SAR (Sandia).

The right-hand (or bottom) image is a low-resolution sarMapper map of the image. The map denotes the material composition (not the physical structure) of the imagery: red denotes possible target, gray denotes shadow, and green denotes vegetation/soil. Note that the T-72 and its shadow are clearly distinguished from the vegetation/soil background.



The montage above shows a Lincoln Lab ADTS X-band, HH-polarized image of the Stockbridge area at the upper left: the image has had an MSTAR T-72 target added to it (highlighted by the yellow bounding box). The lower left shows a low-resolution sarMapper map generated from the SAR image in about two seconds on a laptop. The map legend is to the right: gray denotes shadow, light green denotes grass, dark green denotes trees & shrubs, and red denotes possible targets. Six possible targets are detected in this image, representing less than one percent of the total image area. The T-72 is among the possible targets detected. This result illustrates sarMapper's rapid focus-of-attention capability. Designed to map the material composition of ground cover based on its SAR backscatter signature (rather than its physical structure), sarMapper can distinguish man-made/modified ground cover from natural ground cover in near real-time on a laptop computer. The detected man made/modified ground cover can then be passed on to more computationally intensive ATR algorithms to determine if it represents a target of interest.

 [JBH II Home Page](#)

 [DARPA IU Contract Page](#)

This web page was last updated by ham@sece.cmu.edu on 27-Jul-1997.

© 1997 by John Benjamin Hampshire II; all rights reserved.